

In the claims

1. (currently amended) A method comprising:  
assigning interrupts for a plurality of input/output (I/O) devices among a plurality of nodes of a system based on at least one of: the nodes to which the I/O devices are connected; the nodes at which interrupt service routines for the I/O devices reside; and, processors of the nodes for the nodes having processors, where one or more of the nodes are processorless and/or memoryless and one or more other of the nodes have processors and memory,  
wherein assigning the interrupts for the I/O devices among the nodes of the system comprises, for a given interrupt for a given I/O device, the given I/O device connected to a first node of the nodes of the system, the given I/O device having an interrupt service routine to handle the given interrupt, the interrupt service routine residing at a second node of the nodes of the system, the first node being different than the second node,  
if the first node to which the given I/O device is connected has a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to the first node, even if the second node has a cache, memory, and at least one processor;  
if the first node does not have a cache, memory, and at least one processor, but the second node at which the interrupt service routine of the given I/O device resides does have a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to the second node, even if a third node has memory and at least one processor;  
if both the first node and the second node do not each have a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to [[a]] the third node of the nodes of the system, the third node having memory and at least one processor, the third node being different than the first node and the second node;

for each node having processors, assigning the interrupts for the devices that are performance critical and that have been assigned to the node among the processors of the node in a round-robin manner;

dynamically modifying assignments of the interrupts among the nodes of the system based on actual performance characteristics of the assignments; and,

for each node having processors, dynamically modifying assignments of the interrupts that are performance critical and that have been assigned to the node among the processors of the node based on actual performance characteristics of the assignments.

2. (cancelled)

3. (cancelled)

4. (previously presented) The method of claim 1, wherein dynamically modifying the assignments of the interrupts among the nodes of the system comprises, for each assignment of an interrupt for an I/O device to a node, where the node is that to which the I/O device is connected:

measuring responsiveness of the node in processing the interrupt;

assigning the interrupt to the node at which the interrupt service routine for the I/O device resides;

measuring responsiveness of the node at which the interrupt service routine for the I/O device resides in processing the interrupt; and,

where the responsiveness of the node to which the I/O device is connected is better than the responsiveness of the node at which the interrupt service routine for the I/O device resides, reassigning the interrupt to the node to which the I/O device is connected.

5. (previously presented) The method of claim 4, wherein dynamically modifying the assignments of the interrupts among the nodes of the system comprises, for each assignment of an interrupt for an I/O device to a node, where the node is that at which the interrupt service routine for the I/O device resides:

- measuring responsiveness of the node in processing the interrupt;
- assigning the interrupt to the node to which the I/O device is connected;
- measuring responsiveness of the node to which the I/O device is connected in processing the interrupt; and,

- where the responsiveness of the node at which the interrupt service routine for the I/O device resides is better than the responsiveness of the node to which the I/O device is connected, reassigning the interrupt to the node at which the interrupt service routine for the I/O device resides.

6. (previously presented) The method of claim 1, wherein, for each node of the system having processors, dynamically modifying the assignments of the interrupts that are performance critical and that have been assigned to the node among the processors of the node comprises:

- measuring responsiveness of the processors of the node in processing the interrupts assigned thereto;

- where a differential between a best responsiveness and a worst responsiveness is greater than a threshold,

- reassigning at least one of the interrupts assigned to the processor having the worst responsiveness to the processor having the best responsiveness.

7. (currently amended) A non-uniform memory access (NUMA) system comprising:  
a plurality of nodes;

a plurality of input/output (I/O) devices, each I/O device connected to one of the nodes and having an interrupt; and,

an interrupt-assignor responsive to the I/O devices and the nodes to assign the interrupt for each I/O device to one of the nodes in a performance-optimized manner,

wherein one or more of the nodes are memoryless and/or processorless, and one or more other of the nodes have memory and processors, and

wherein the interrupt-assignor assigns a given interrupt for a given I/O device of the plurality of I/O devices, the given I/O device connected to a first node of the nodes of the system, the given I/O device having an interrupt service routine to handle the given interrupt, the interrupt service routine residing at a second node of the nodes of the system, the first node being different than the second node, the interrupt-assignor performing:

if the first node to which the given I/O device is connected has a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to the first node, even if the second node has a cache, memory, and at least one processor;

if the first node does not have a cache, memory, and at least one processor, but the second node at which the interrupt service routine of the given I/O device resides does have a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to the second node, even if a third node has memory and at least one processor;

if both the first node and the second node do not each have a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to ~~[[a]]~~ the third node of the nodes of the system, the third node having memory and at least one processor, the third node being different than the first node and the second node.

8. (previously presented) The system of claim 7, wherein the memory of each node that has memory is local to the node and remote to all other of the nodes, and the interrupt-assignor is to

assign the interrupt for each I/O device to one of the nodes that has memory and at least one processor.

9. (original) The system of claim 8, wherein at least one of the I/O devices are performance critical, the interrupt-assignor further to assign the interrupt for each I/O device that is performance critical among the at least one processor of the node to which the interrupt has been assigned in a round-robin manner.

10. (previously presented) The system of claim 7, wherein, for each node that has processors, the interrupt-assignor is further to dynamically modify assignments of the interrupts that are performance critical among the at least one processor of the node based on actual performance characteristics of the assignments.

11. (cancelled)

12. (original) The system of claim 7, wherein the interrupt-assignor is further to dynamically modify assignments of the interrupts among the plurality of nodes based on actual performance characteristics of the assignments.

13. (original) The system of claim 7, wherein the interrupt-assignor is to give primary preference in assigning the interrupt for each I/O device to the node to which the I/O device is connected where the node to which the I/O device is connected has a cache, memory, and at least one processor.

14. (original) The system of claim 13, wherein each I/O device further has an interrupt service routine residing at one of the plurality of nodes, and the interrupt-assignor is to give secondary

preference in assigning the interrupt for each I/O device to the node at which the interrupt service routine of the I/O device resides where the node at which the interrupt service routine of the I/O device resides has a cache, memory, and at least one processor.

15. (original) The system of claim 7, wherein the interrupt-assignor resides within one of the plurality of nodes.

16. (currently amended) An article of manufacture comprising:

a computer-readable storage medium; and,

means in the medium, executed by a processor, for assigning interrupts for a plurality of input/output (I/O) devices among a plurality of nodes of a system based on at least one factor selected from the set consisting of: the nodes to which the I/O devices are connected; and the nodes at which interrupt service routines for the I/O devices reside, where one or more of the nodes are processorless and/or memoryless and one or more other of the nodes have processors and memory,

wherein the means assigns the interrupts by for a given interrupt for a given I/O device, the given I/O device connected to a first node of the nodes of the system, the given I/O device having an interrupt service routine to handle the given interrupt, the interrupt service routine residing at a second node of the nodes of the system, the first node being different than the second node,

if the first node to which the given I/O device is connected has a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to the first node, even if the second node has a cache, memory, and at least one processor;

if the first node does not have a cache, memory, and at least one processor, but the second node at which the interrupt service routine of the given I/O device resides does have a cache, memory, and at least one processor, then assigning the given interrupt for the given

I/O device to the second node, even if a third node has memory and at least one processor;

if both the first node and the second node do not each have a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to [[a]] the third node of the nodes of the system, the third node having memory and at least one processor, the third node being different than the first node and the second node.

17. (original) The article of claim 16, wherein the means is for assigning the interrupts among the plurality of nodes further based on whether the nodes have processors and memories.

18. (previously presented) The article of claim 16, wherein the means, for each node having processors, is further for assigning the interrupts for the devices that are performance critical and that have been assigned to the node among the processors of the node in a round-robin manner.

19. (previously presented) The article of claim 18, wherein the means, is further for dynamically modifying assignments of the interrupts among the nodes based on actual performance characteristics of the assignments, and, for each node having processors, for dynamically modifying assignments of the interrupts that are performance critical and that have been assigned to the node among the processors of the node based on actual performance characteristics of the assignments.

20. (cancelled)

21. (currently amended) A computer-readable storage medium comprising:  
an interrupt-assignor to assign interrupts for a plurality of input/output (I/O) devices among a plurality of nodes of a system based on at least one factor selected from the set consisting of: the nodes to which the I/O devices are connected; and the nodes at which interrupt

service routines for the I/O devices reside, where one or more of the nodes are processorless and/or memoryless and one or more other of the nodes have processors and memory,

wherein the interrupt-assignor assigns a given interrupt for a given I/O device of the plurality of I/O devices, the given I/O device connected to a first node of the nodes of the system, the given I/O device having an interrupt service routine to handle the given interrupt, the interrupt service routine residing at a second node of the nodes of the system, the first node being different than the second node, the interrupt-assignor performing:

if the first node to which the given I/O device is connected has a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to the first node, even if the second node has a cache, memory, and at least one processor;

if the first node does not have a cache, memory, and at least one processor, but the second node at which the interrupt service routine of the given I/O device resides does have a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to the second node, even if a third node has memory and at least one processor;

if both the first node and the second node do not each have a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to ~~[[a]]~~ the third node of the nodes of the system, the third node having memory and at least one processor, the third node being different than the first node and the second node.

22. (previously presented) The medium of claim 21, wherein the interrupt-assignor is to assign the interrupts among the plurality of nodes further based on whether the nodes have processors and memories.

23. (previously presented) The medium of claim 21, wherein the interrupt-assignor is to assign, for each node, the interrupts for the devices that are performance critical and that have been assigned to the node among the processors of the node in a round-robin manner.



24. (previously presented) The medium of claim 23, wherein the interrupt-assignor is to dynamically modify assignments of the interrupts among the nodes based on actual performance characteristics of the assignments, and, for each node, to dynamically modify assignments of the interrupts that are performance critical and that have been assigned to the node among the processors of the node based on actual performance characteristics of the assignments.

25. (currently amended) A method comprising:

assigning interrupts for a plurality of input/output (I/O) devices among a plurality of nodes based on at least one factor selected from the set consisting of: the nodes to which the I/O devices are connected; and the nodes at which interrupt service routines for the I/O devices reside, where one or more of the nodes are processorless and/or memoryless and one or more other of the nodes have processors and memory,

wherein assigning the interrupts for the I/O devices among the nodes of the system comprises, for a given interrupt for a given I/O device, the given I/O device connected to a first node of the nodes of the system, the given I/O device having an interrupt service routine to handle the given interrupt, the interrupt service routine residing at a second node of the nodes of the system, the first node being different than the second node,

if the first node to which the given I/O device is connected has a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to the first node, even if the second node has a cache, memory, and at least one processor;

if the first node does not have a cache, memory, and at least one processor, but the second node at which the interrupt service routine of the given I/O device resides does have a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to the second node, even if a third node has memory and at least one processor;

if both the first node and the second node do not each have a cache, memory,

and at least one processor, then assigning the given interrupt for the given I/O device to [[a]] the third node of the nodes of the system, the third node having memory and at least one processor, the third node being different than the first node and the second node;

for each node of the system having processors, assigning the interrupts for the devices that are performance critical and that have been assigned to the node among the processors of the node in a round-robin manner;

dynamically modifying assignments of the interrupts among the nodes of the system based on actual performance characteristics of the assignments; and,

for each node of the system having processors, dynamically modifying assignments of the interrupts that are performance critical and that have been assigned to the node among the processors of the node based on actual performance characteristics of the assignments.

26. (cancelled)

27. (cancelled)

28. (original) The method of claim 25, wherein dynamically modifying the assignments of the interrupts among the nodes of the system comprises, for each assignment of an interrupt for an I/O device to a node, where the node is that to which the I/O device is connected:

measuring responsiveness of the node in processing the interrupt;

assigning the interrupt to the node at which the interrupt service routine for the I/O device resides;

measuring responsiveness of the node at which the interrupt service routine for the I/O device resides in processing the interrupt; and,

where the responsiveness of the node to which the I/O device is connected is better than the responsiveness of the node at which the interrupt service routine for the I/O device resides, reassigning the interrupt to the node to which the I/O device is connected.

29. (original) The method of claim 28, wherein dynamically modifying the assignments of the interrupts among the nodes of the system comprises, for each assignment of an interrupt for an I/O device to a node, where the node is that at which the interrupt service routine for the I/O device resides:

- measuring responsiveness of the node in processing the interrupt;
- assigning the interrupt to the node to which the I/O device is connected;
- measuring responsiveness of the node to which the I/O device is connected in processing the interrupt; and,

where the responsiveness of the node at which the interrupt service routine for the I/O device resides is better than the responsiveness of the node to which the I/O device is connected, reassigning the interrupt to the node at which the interrupt service routine for the I/O device resides.

30. (previously presented) The method of claim 25, wherein, for each node of the system having memory, dynamically modifying the assignments of the interrupts that are performance critical and that have been assigned to the node among the processors of the node comprises:

- measuring responsiveness of the processors of the node in processing the interrupts assigned thereto;

where a differential between a best responsiveness and a worst responsiveness is greater than a threshold,

reassigning at least one of the interrupts assigned to the processor having the worst responsiveness to the processor having the best responsiveness.